L Number	Hits	Search Text	DB	Time stamp
-	0	(positon adj information) and (identification adj information) and hierarchical	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 08:13
-	2680	directory adj service	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 08:40
-	942	(directory adj service) and hierarchic$5	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 08:40
-	365	((directory adj service) and hierarchic$5) and synchroni$3	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 08:40
-	138	(((directory adj service) and hierarchic$5) and synchroni$3) and (@rlad<=19990630 or @ad<=19990630)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 12:46
-	2	"6466932"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 09:00
-	1	6539381.URPN.	USPAT	2003/09/10 09:19
-	8	("5596574" | "5765171" | "5787247" | "5832487" | "6021118" | "6065017" | "6098078" | "6233623").PN.	USPAT	2003/09/10 09:19
-	3190	707/10.ccls.	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 12:45
-	93	707/10.ccls. and LDAP	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 12:46
-	52	(707/10.ccls. and LDAP) and hierarchical	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 12:46
-	27	((707/10.ccls. and LDAP) and hierarchical) and (@rlad<=19990630 or @ad<=19990630)	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 12:47
-	24	("5291583" | "5317568" | "5317742" | "5331642" | "5471617" | "5473608" | "5491694" | "5491796" | "5509123" | "5517622" | "5522042" | "5541911" | "5548796" | "5559958" | "5561769" | "5566160" | "5572195" | "5586304" | "5588143" | "5721909" | "5764977" | "5799153" | "5835757" | "5878434").PN.	USPAT	2003/09/10 14:36
-	19	"6052724"	USPAT; US-PGPUB; EPO; JPO; IBM_TDB	2003/09/10 14:54

Terms used **hierach directory synchroniz** Found **1,782** of **121,259**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 SoftFLASH: analyzing the performance of clustered distributed virtual shared memory

Andrew Erlichson, Neal Nuckolls, Greg Chesson, John Hennessy

September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, citings, index terms

One potentially attractive way to build large-scale shared-memory machines is to use small-scale to medium-scale shared-memory machines as clusters that are interconnected with an off-the-shelf network. To create a shared-memory programming environment across the clusters, it is possible to use a virtual shared-memory software layer. Because of the low latency and high bandwidth of the interconnect available within each cluster, there are clear advantages in making the clusters as large as possi ...

2 A computer aided instruction (CAI) course for learning Ada

Richard C. Felsinger

March 1986 **Proceedings of the third annual Washington Ada symposium on Ada: Ada use in focus : practical lessons in perspective**

Full text available: pdf(1.42 MB) Additional Information: full citation, references

3 Structured specification of a Security Kernel

K. G. Walter, S. I. Schaen, W. F. Ogden, W. C. Rounds, D. G. Shumway, D. D. Schaeffer, K. J. Biba, F. T. Bradshaw, S. R. Ames, J. M. Gilligan

April 1975 **ACM SIGPLAN Notices , Proceedings of the international conference on Reliable software**, Volume 10 Issue 6

Full text available: pdf(745.99 KB) Additional Information: full citation, abstract, references, citings, index terms

Certifying an entire operating system to be reliable is too large a task to be practicable. Instead, we are designing a Security Kernel which will provide information security. The kernel's job is to monitor information flow in order to prevent compromise of security. Sound design is encouraged by using a technique called Structured Specification, in which successively more detailed models of the Security Kernel are developed. The initial model, M ...

Keywords: Compromise security, Governmental security, Model, Security kernel, Security system, Structured programming, Structured specification

4 Disk cache—miss ratio analysis and design considerations
Alan J. Smith
August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Full text available: pdf(3.13 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The current trend of computer system technology is toward CPUs with rapidly increasing processing power and toward disk drives of rapidly increasing density, but with disk performance increasing very slowly if at all. The implication of these trends is that at some point the processing power of computer systems will be limited by the throughput of the input/output (I/O) system. A solution to this problem, which is described and evaluated in this paper, is disk cache

5 Synchronizing shared abstract types
Peter M. Schwarz, Alfred Z. Spector
August 1984 **ACM Transactions on Computer Systems (TOCS)**, Volume 2 Issue 3

Full text available: pdf(1.93 MB) Additional Information: full citation, references, citings, index terms

Keywords: dependencies, locking, transaction serializability

6 An algebraic approach to file synchronization
Norman Ramsey, El&ohuml;d Csirmaz
September 2001 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 8th European software engineering conference held jointly with 9th ACM SIGSOFT international symposium on Foundations of software engineering**, Volume 26 Issue 5

Full text available: pdf(301.78 KB) Additional Information: full citation, abstract, references, index terms

A *file synchronizer* restores consistency after multiple replicas of a filesystem have been changed independently. We present an algebra for reasoning about operations on filesystems and show that it is sound and complete with respect to a simple model. The algebra enables us to specify a file-synchronization algorithm that can be combined with several different conflict-resolution policies. By contrast, previous work builds the conflict-resolution policy into the specification, or worse, ...

7 One user, one password: integrating unix accounts and active directory
David J. Blezard, Jerry Marceau
November 2002 **Proceeding of the 30th annual ACM SIGUCCS fall conference on User services conference**

Full text available: pdf(141.57 KB) Additional Information: full citation, abstract, references, index terms

The University of New Hampshire has a history of using centralized Unix accounts to authenticate user access to computers in the public Student Computing Clusters. The advent of Windows 2000 meant that changes would be necessary to support the Active Directory architecture underlying Windows 2000 authentication and authorization. Given limited resources, manually maintaining Active Directory accounts for over 12000 students is an impossibility. A new system was needed to automatically generate a ...

Keywords: Active Directory, Unix, Windows 2000, accounts, authentication, integration

8 Multiprocessor cache synchronization: issues, innovations, evolution
P. Bitar, A. M. Despain

June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2

Full text available: pdf(981.69 KB) Additional Information: full citation, abstract, references, citings, index terms

Many options are possible in a cache synchronization (or consistency) scheme for a broadcast system. We clarify basic concepts, analyze the handling of shared data, and then describe a protocol that we are currently exploring. Finally, we analyze the evolution of options that have been proposed under write-in (or write-back) policy. We show how our protocol extends this evolution with new methods for efficient busy-wait locking, waiting, and unlocking. The ...

9 A weighted voting algorithm for replicated directories

Joshua J. Bloch, Dean S. Daniels, Alfred Z. Spector

October 1987 **Journal of the ACM (JACM)**, Volume 34 Issue 4

Full text available: pdf(4.12 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Weighted voting is used as the basis for a replication technique for directories. This technique affords arbitrarily high data availability as well as high concurrency. Efficient algorithms are presented for all of the standard directory operations. A structural property of the replicated directory that permits the construction of an efficient algorithm for deletion is proven. Simulation results are presented and the system is modeled and analyzed. The analysis agrees well with the simulati ...

10 An empirical evaluation of two memory-efficient directory methods

Brian W. O'Krafka, A. Richard Newton

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents an empirical evaluation of two memory-efficient directory methods for maintaining coherent caches in large shared memory multiprocessors. Both directory methods are modifications of a scheme proposed by Censier and Feautrier [5] that does not rely on a specific interconnection network and can be readily distributed across interleaved main memory. The schemes considered here overcome the large amount of memory required for tags in the original scheme in two different ways ...

11 What is a file synchronizer?

S. Balasubramaniam, Benjamin C. Pierce

October 1998 **Proceedings of the 4th annual ACM/IEEE international conference on Mobile computing and networking**

Full text available: pdf(1.21 MB) Additional Information: full citation, references, citings, index terms

12 LimitLESS directories: A scalable cache coherence scheme

David Chaiken, John Kubiatowicz, Anant Agarwal

April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the fourth international conference on Architectural support for programming languages and operating systems**, Volume 19 Issue 2

Full text available: pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

13 The directory-based cache coherence protocol for the DASH multiprocessor

Daniel Lenoski, James Laud[redacted] Kourosh Gharachorloo, Anoop Gupta, [redacted]n Hennessy
May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, citings, index terms

DASH is a scalable shared-memory multiprocessor currently being developed at Stanford's Computer Systems Laboratory. The architecture consists of powerful processing nodes, each with a portion of the shared-memory, connected to a scalable interconnection network. A key feature of DASH is its distributed directory-based cache coherence protocol. Unlike traditional snoopy coherence protocols, the DASH protocol does not rely on broadcast; instead it uses point-to-point messages sent between th ...

14 Adaptive backoff synchronization techniques

A. Agarwal, M. Cherian
April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references, citings, index terms

Shared-memory multiprocessors commonly use shared variables for synchronization. Our simulations of real parallel applications show that large-scale cache-coherent multiprocessors suffer significant amounts of invalidation traffic due to synchronization. Large multiprocessors that do not cache synchronization variables are often more severely impacted. If this synchronization traffic is not reduced or managed adequately, synchronization references can cause severe congestion in the network. ...

15 Mobility and Wireless Access: Personalized pocket directories for mobile devices

Doron Cohen, Michael Herscovici, Yael Petruschka, Yoëlle S. Maarek, Aya Soffer
May 2002 **Proceedings of the eleventh international conference on World Wide Web**

Full text available: pdf(529.16 KB) Additional Information: full citation, abstract, references, index terms

In spite of the increase in the availability of mobile devices in the last few years, Web information is not yet as accessible from PDAs or WAP phones as it is from the desktop. In this paper, we propose a solution for supporting one of the most popular information discovery mechanisms, namely Web directory navigation, from mobile devices. Our proposed solution consists of caching enough information on the device itself in order to conduct most of the navigation actions locally (with subsecond r ...

Keywords: hierarchical browsers, mobile devices, mobile search, personalization

16 An algorithm, for replicated directories

Dean Daniels, Alfred Z. Spector
August 1983 **Proceedings of the second annual ACM symposium on Principles of distributed computing**

Full text available: pdf(825.73 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes a replication algorithm for directory objects based upon Gifford's weighted voting for files. The algorithm associates version number with each possible key on every replica and thereby resolves an ambiguity that arises when directory entries are not stored in every replica. The range of keys associated with a version number changes dynamically; but in all instances, a separate version number is associated with each entry stored on every replica. The algorithm exhibits ...

Keywords: Availability, Replicated data, Transaction- based systems

17 Algorithms for scalable synchronization on shared-memory multiprocessors
John M. Mellor-Crummey, Michael L. Scott
February 1991 **ACM Transactions on Computer Systems (TOCS)**, Volume 9 Issue 1

Full text available: pdf(3.07 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Busy-wait techniques are heavily used for mutual exclusion and barrier synchronization in shared-memory parallel programs. Unfortunately, typical implementations of busy-waiting tend to produce large amounts of memory and interconnect contention, introducing performance bottlenecks that become markedly more pronounced as applications scale. We argue that this problem is not fundamental, and that one can in fact construct busy-wait synchronization algorithms that induce no memory or interc ...

18 Parallel architectures: Inferential queueing and speculative push for reducing critical communication latencies
Ravi Rajwar, Alain Kägi, James R. Goodman
June 2003 **Proceedings of the 17th annual international conference on Supercomputing**

Full text available: pdf(568.93 KB) Additional Information: full citation, abstract, references, index terms

Communication latencies within critical sections constitute a major bottleneck in some classes of emerging parallel workloads. In this paper, we argue for the use of Inferentially Queued Locks (IQLs) [31], not just for efficient synchronization but also for reducing communication latencies, and we propose a novel mechanism, Speculative Push (SP), aimed at reducing these communication latencies. With IQLs, the processor infers the existence, and limits, of a critical section from the use of synch ...

Keywords: data forwarding, inferential queueing, synchronization

19 Synchronization and recovery of actions
J. E. Allchin, M. S. McKendry
August 1983 **Proceedings of the second annual ACM symposium on Principles of distributed computing**

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, citings, index terms

We introduce an approach to robust computation in distributed systems. This approach is the foundation for reliability in the Clouds decentralized operating system. It is based on atomic actions operating on instances of abstract data types (objects). We present an event-based model of computation in which scheduling of responses to operation invocations is controlled by objects. We discuss an integrated strategy for synchronization and recovery which uses rela ...

20 Architectural primitives for a scalable shared memory multiprocessor
Joonwon Lee, Umakishore Ramachandran
June 1991 **Proceedings of the third annual ACM symposium on Parallel algorithms and architectures**

Full text available: pdf(1.27 MB) Additional Information: full citation, references, citings, index terms